

Statistical depth meets computational geometry: a short survey

Peter J. Rousseeuw and Mia Hubert

August 15, 2015

Abstract

During the past two decades there has been a lot of interest in developing statistical depth notions that generalize the univariate concept of ranking to multivariate data. The notion of depth has also been extended to regression models and functional data. However, computing such depth functions as well as their contours and deepest points is not trivial. Techniques of computational geometry appear to be well-suited for the development of such algorithms. Both the statistical and the computational geometry communities have done much work in this direction, often in close collaboration. We give a short review of this work, focusing mainly on depth and multivariate medians, and end by listing some other areas of statistics where computational geometry has been of great help in constructing efficient algorithms.

1 Depth notions

A data set consisting of n univariate points is usually ranked in ascending or descending order. Univariate order statistics (i.e., the ' k th smallest value out of n ') and derived quantities have been studied extensively. The median is defined as the order statistic of rank $(n + 1)/2$ when n is odd, and as the average of the order statistics of ranks $n/2$ and $(n + 2)/2$ when n is even. The median and any other order statistic of a univariate data set can be computed in $O(n)$ time. Generalization to higher dimensions is, however, not straightforward.

Alternatively, univariate points may be ranked from the outside inward by assigning the most extreme data points depth 1, the second smallest and second largest data points depth 2, etc. The deepest point then equals the usual median of the sample. The advantage of this type of ranking is that it can be extended to higher dimensions more easily. This section gives an overview of several possible generalizations of depth and the median to multivariate settings. Surveys of statistical applications of multivariate data depth may be found in [LPS99], [ZS00], and [Mos13].

1.1 Halfspace location depth

Let $X_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ be a finite set of data points in \mathbb{R}^d . The *Tukey depth* or *halfspace depth* (introduced by [Tuk75] and further developed by [DG92]) of any point $\boldsymbol{\theta}$ in \mathbb{R}^d (not necessarily a data point) determines how central the point is inside the data cloud. The halfspace depth of $\boldsymbol{\theta}$ is defined as the minimal number of data points in any closed halfspace determined by a hyperplane through $\boldsymbol{\theta}$:

$$hdepth(\boldsymbol{\theta}; X_n) = \min_{\|\mathbf{u}\|=1} \#\{i; \mathbf{u}^\top \mathbf{x}_i \geq \mathbf{u}^\top \boldsymbol{\theta}\}.$$

Thus, a point lying outside the convex hull of X_n has depth 0, and any data point has depth at least 1. Figure 1 illustrates this definition for $d = 2$.

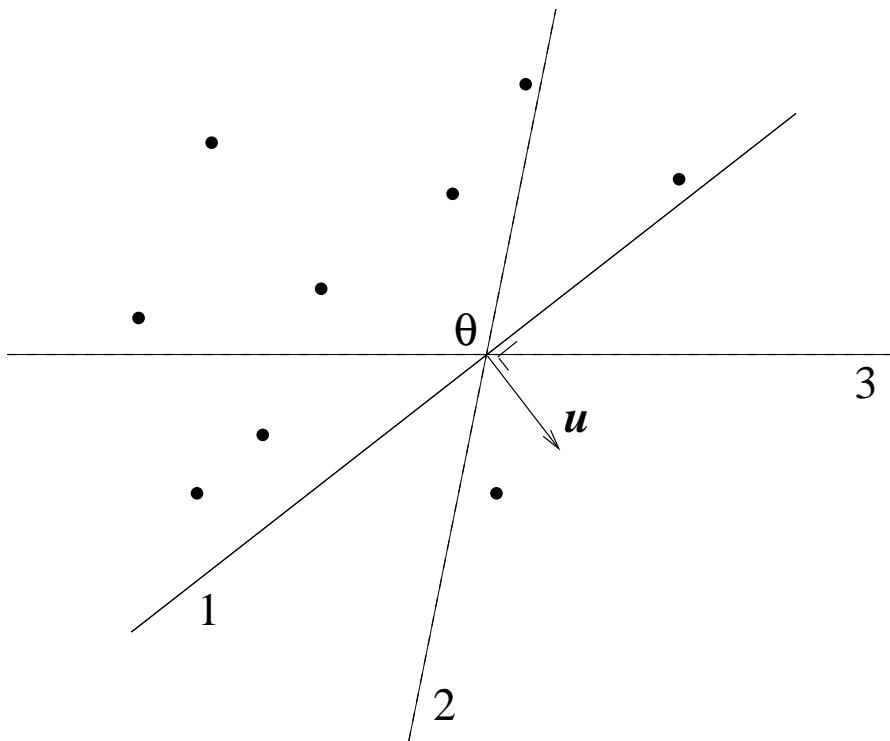


Figure 1: Illustration of the bivariate halfspace depth. Here $\boldsymbol{\theta}$ (which is not a data point itself) has depth 1 because the halfspace determined by \mathbf{u} contains only one data point.

The set of all points with depth $\geq k$ is called the k th depth region D_k . The halfspace depth regions form a sequence of nested polyhedra. Each D_k is the intersection of all halfspaces containing at least $n - k + 1$ data points. Moreover, every data point must be a vertex of one or more depth regions. The point with maximal halfspace depth is called the *Tukey median*. When the innermost depth region is larger than a singleton, the Tukey median is defined as its centroid. This makes the Tukey median unique by construction.

Note that the depth regions give an indication of the shape of the data cloud. Based on this idea one can construct the *bagplot* [RRT99], a bivariate version of the univariate boxplot. Figure 2 shows such a bagplot. The cross in the white disk is the Tukey median. The dark area is an interpolation between two subsequent depth regions, and contains 50% of the data. This area (the “bag”) gives an idea of the shape of the majority of the data cloud. Inflating the bag by a factor of 3 relative to the Tukey median yields the “fence” (not shown), and data points outside the fence are called outliers and marked by stars. Finally, the light gray area is the convex hull of the non-outlying data points.

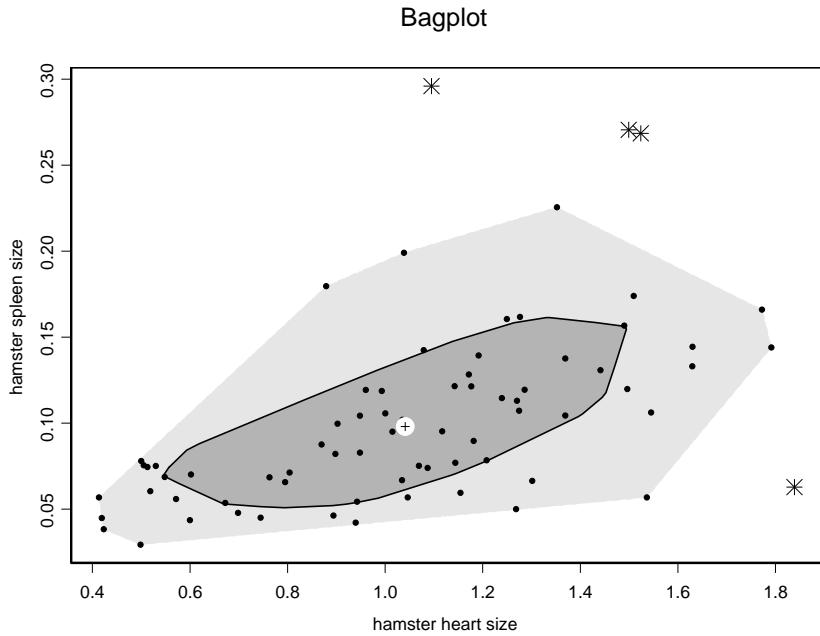


Figure 2: Bagplot of the heart and spleen size of 73 hamsters.

More generally, in the multivariate case one can define the *bagdistance* [HRS15b] of a point \mathbf{x} relative to the Tukey median and the bag. Assume that the Tukey median lies in the interior of the bag, not on its boundary (this excludes degenerate cases). Then the bagdistance is the smallest real number λ such that the bag inflated (or deflated) by λ around the Tukey median contains the point \mathbf{x} . When the Tukey median equals $\mathbf{0}$, it is shown in [HRS15b] that the bagdistance satisfies all axioms of a norm except that $\|a\mathbf{x}\| = |a|\|\mathbf{x}\|$ only needs to hold when $a \geq 0$. The bagdistance is used for outlier detection [HRS15a] and statistical classification [HRS15b].

An often used criterion to judge the robustness of an estimator is its *breakdown value*. The breakdown value is the smallest fraction of data points that we need

to replace in order to move the estimator of the contaminated data set arbitrarily far away. The classical mean of a data set has breakdown value zero since we can move it anywhere by moving one observation. Note that for any estimator which is equivariant for translation (which is required to call it a location estimator) the breakdown value can be at most $1/2$. (If we replace half of the points by a far-away translation image of the remaining half, the estimator cannot distinguish which were the original data.)

The Tukey depth and the corresponding median have good statistical properties. The Tukey median T^* is a location estimator with breakdown value $\varepsilon_n(T^*; X_n) \geq 1/(d+1)$ for any data set in general position. This means that it remains in a predetermined bounded region unless $n/(d+1)$ or more data points are moved. At an elliptically symmetric distribution the breakdown value becomes $1/3$ for large n , irrespective of d . Moreover, the halfspace depth is invariant under all nonsingular affine transformations of the data, making the Tukey median affine equivariant. Since data transformations such as rotation and rescaling are very common in statistics, this is an important property. The statistical asymptotics of the Tukey median have been studied in [BH99].

The need for fast algorithms for the halfspace depth has only grown over the years, since it is currently being applied to a variety of settings such as nonparametric classification [LCL12]. A related development is the fast growing field of functional data analysis, where the data are functions on a univariate interval (e.g. time or wavelength) or on a rectangle (e.g. surfaces, images). Often the function values are themselves multivariate. One can then define the depth of a curve (surface) by integrating the depth over all points as in [Cla14]. This functional depth can again be used for outlier detection and classification [HRS15a, HRS15b], but it requires computing depths in many multivariate data sets instead of just one.

Remark: centerpoints. There is a close relationship between the Tukey depth and centerpoints, which have been long studied in computational geometry. In fact, Tukey depth extends the notion of centerpoint. A *centerpoint* is any point with halfspace depth $\geq \lceil n/(d+1) \rceil$. A consequence of Helly's theorem is that there always exists at least one centerpoint, so the depth of the Tukey median cannot be less than $\lceil n/(d+1) \rceil$.

1.2 Other location depth notions

1. Simplicial depth ([Liu90]). The depth of $\boldsymbol{\theta}$ equals the number of simplices formed by $d+1$ data points that contain $\boldsymbol{\theta}$. Formally,

$$sdepth(\boldsymbol{\theta}; X_n) = \#\{(i_1, \dots, i_{d+1}); \boldsymbol{\theta} \in S[\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_{d+1}}]\}.$$

The simplicial median is affine equivariant with a breakdown value bounded above by $1/(d+2)$. Unlike halfspace depth, its depth regions need not be convex.

2. Oja depth ([Oja83]). This is also called simplicial volume depth:

$$odepth(\boldsymbol{\theta}; X_n) = \left(1 + \sum_{(i_1, \dots, i_d)} \{volume\ S[\boldsymbol{\theta}, \mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_d}]\}\right)^{-1}.$$

The corresponding median is also affine equivariant, but has zero breakdown value.

3. Projection depth. We first define the *outlyingness* ([DG92]) of any point $\boldsymbol{\theta}$ relative to the data set X_n as

$$O(\boldsymbol{\theta}; X_n) = \max_{\|\mathbf{u}\|=1} \frac{|\mathbf{u}^\tau \boldsymbol{\theta} - \text{median}_i \{\mathbf{u}^\tau \mathbf{x}_i\}|}{\text{MAD}_i \{\mathbf{u}^\tau \mathbf{x}_i\}},$$

where the median absolute deviation (MAD) of a univariate data set $\{y_1, \dots, y_n\}$ is the statistic $\text{MAD}_i \{y_i\} = \text{median}_i |y_i - \text{median}_j \{y_j\}|$. The outlyingness is small for centrally located points and increases if we move toward the boundary of the data cloud. Instead of the median and the MAD, also another pair (T, S) of a location and scatter estimate may be chosen. This leads to different notions of projection depth, all defined as

$$pdepth(\boldsymbol{\theta}; X_n) = (1 + O(\boldsymbol{\theta}; X_n))^{-1}.$$

General projection depth is studied in [Zuo03]. When using the median and the MAD, the projection depth has breakdown value $1/2$ and is affine equivariant. Its depth regions are convex.

4. Spatial depth ([Ser02]). Spatial depth is related to multivariate quantiles proposed in [Cha96]:

$$spdepth(\boldsymbol{\theta}; X_n) = 1 - \left\| \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{x}_i - \boldsymbol{\theta}}{\|\mathbf{x}_i - \boldsymbol{\theta}\|} \right\|$$

The spatial median is also called the L^1 median ([Gow74]). It has breakdown value $1/2$, but is not affine equivariant (it is only equivariant with respect to translations, multiplication by a scalar factor, and orthogonal transformations). For a recent survey on the computation of the spatial median see [FFC12].

A comparison of the main properties of the different location depth medians is given in Table 1.

Table 1: Comparison of several location depth medians

MEDIAN	BREAKDOWN VALUE	AFFINE EQUIVARIANCE
Tukey	worst-case $1/(d+1)$ typically $1/3$	yes
Simplicial	$\leq 1/(d+2)$	yes
Oja	$2/n \approx 0$	yes
Projection	$1/2$	yes
Spatial	$1/2$	no

1.3 Arrangement and regression depth

Following [RH99b] we now define the depth of a point relative to an arrangement of hyperplanes. A point $\boldsymbol{\theta}$ is said to have zero arrangement depth if there exists a ray $\{\boldsymbol{\theta} + \lambda\mathbf{u}; \lambda \geq 0\}$ that does not cross any of the hyperplanes h_i in the arrangement. (A hyperplane parallel to the ray is counted as intersecting at infinity.) The arrangement depth of any point $\boldsymbol{\theta}$ is defined as the minimum number of hyperplanes intersected by any ray from $\boldsymbol{\theta}$. Figure 3 shows an arrangement of lines. In this plot, the points $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$ have arrangement depth 0 and the point $\boldsymbol{\xi}$ has arrangement depth 2. The arrangement depth is always constant on open cells and on cell edges. It was shown ([RH99b]) that any arrangement of lines in the plane encloses a point with arrangement depth at least $\lceil n/3 \rceil$, giving rise to a new type of ‘‘centerpoints.’’

This notion of depth was originally defined ([RH99]) in the dual, as the depth of a regression hyperplane $H_{\boldsymbol{\theta}}$ relative to a point configuration of the form $Z_n = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ in \mathbb{R}^{d+1} . Regression depth ranks hyperplanes according to how well they fit the data in a regression model, with \mathbf{x} containing the predictor variables and y the response. A vertical hyperplane (given by $\mathbf{a}^\top \mathbf{x} = \text{constant}$), which cannot be used to predict future response values, is called a ‘‘nonfit’’ and assigned regression depth 0. The regression depth of a general hyperplane $H_{\boldsymbol{\theta}}$ is found by rotating $H_{\boldsymbol{\theta}}$ in a continuous movement until it becomes vertical. The minimum number of data points that is passed in such a rotation is called the regression depth of $H_{\boldsymbol{\theta}}$. Figure 4 is the dual representation of Figure 3. (For instance, the line $\boldsymbol{\theta}$ has slope θ_1 and intercept θ_2 and corresponds to the point (θ_1, θ_2) in Figure 3.) The lines $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$ have regression depth 0, whereas the line $\boldsymbol{\xi}$ has regression depth 2.

In statistics one is interested in the *deepest fit* or regression depth median, because this is a line (hyperplane) about which the data are well-balanced. The statistical properties of regression depth and the deepest fit are very similar to those of the Tukey depth and median. The bounds on the maximal depth are almost the same. Moreover, for both depth notions the value of the maximal depth can be used to characterize the symmetry of the distribution ([RS04]). The

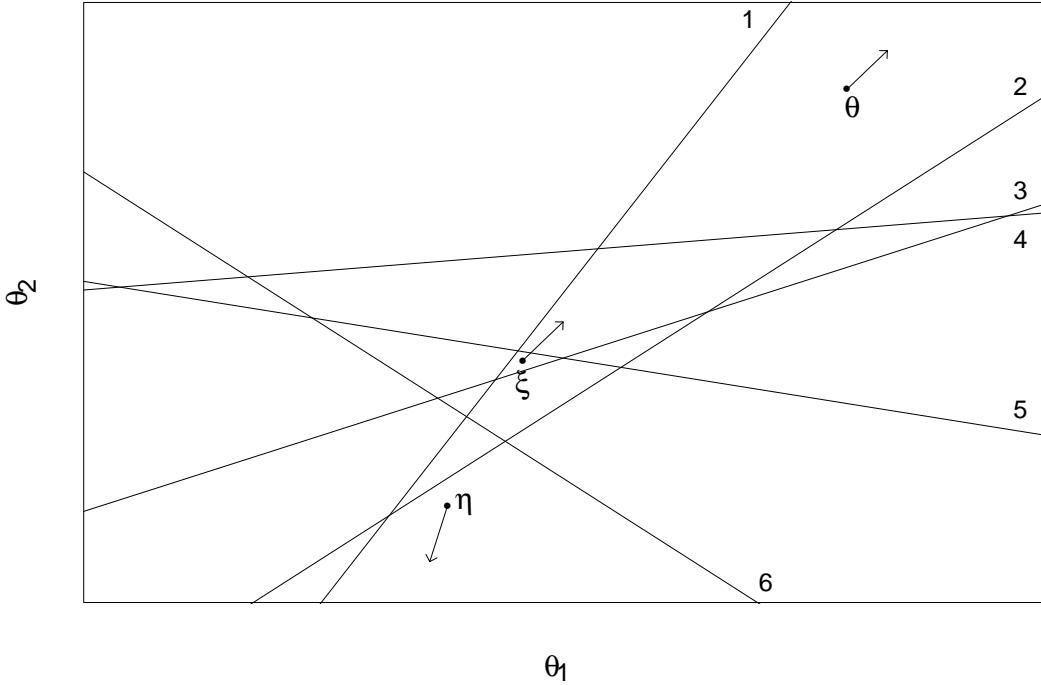


Figure 3: Example of arrangement depth. In this arrangement of lines, the points θ and η have arrangement depth 0, whereas ξ has arrangement depth 2. (See Figure 4 for the dual plot.)

breakdown value of the deepest fit is at least $1/(d+1)$ and under linearity of the conditional median of y given \mathbf{x} it converges to $1/3$. In the next section, we will see that the optimal complexities for computing the depth and the median are also comparable to those for halfspace depth. For a detailed comparison of the properties of halfspace and regression depth, see [HRV01].

The arrangement depth region D_k is defined in the primal, as the set of points with arrangement depth at least k . Contrary to the Tukey depth, these depth regions need not be convex. But nevertheless it was proved that there always exists a point with arrangement depth at least $\lceil n/(d+1) \rceil$ ([ABE⁺00]). An analysis-based proof was given in [Miz02].

Remark: arrangement levels. Arrangement depth is undirected (isotropic) in the sense that it is defined as a minimum over all possible directions. If we restrict ourselves to vertical directions \mathbf{u} (i.e., up or down), we obtain the usual levels of the arrangement known in combinatorial geometry. The absence of preferential directions makes arrangement depth invariant under affine transformations.

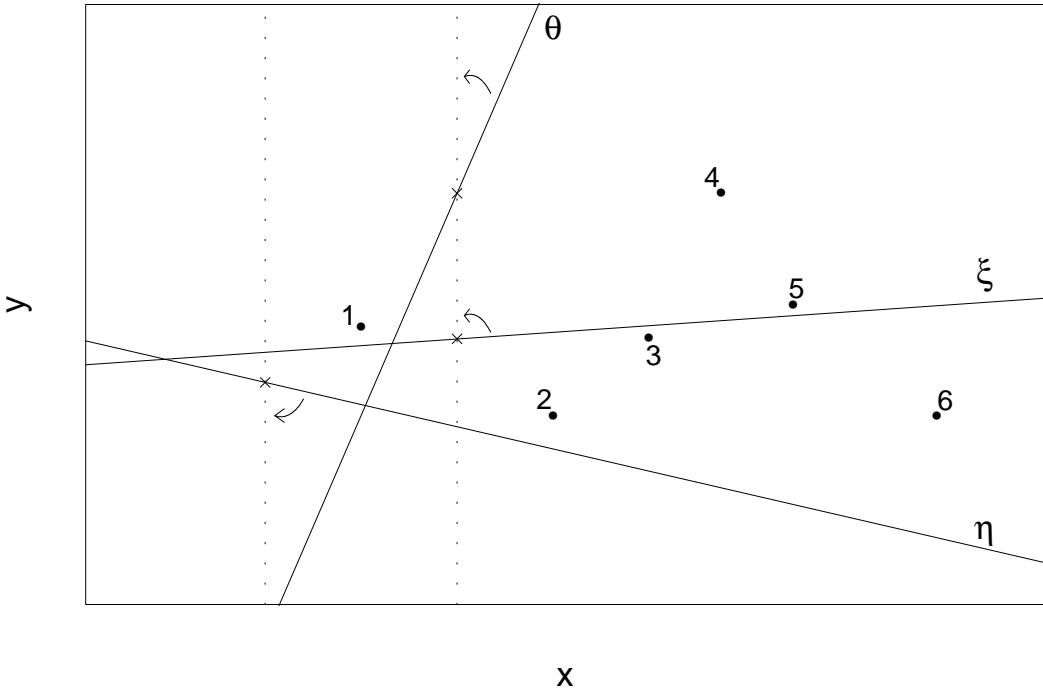


Figure 4: Example of the regression depth of a line in a bivariate configuration of points. The lines θ and η have regression depth 0, whereas the line ξ has regression depth 2. (This is the dual of Figure 3.)

2 Computing depth

Although the definitions of depth are intuitive, the computational aspects can be quite challenging. The calculation of depth regions and medians is computationally intensive, especially for large data sets in higher dimensions. In statistical practice such data are quite common, and therefore reliable and efficient algorithms are needed. For the bivariate case several algorithms have been developed early on. The computational aspects of depth in higher dimensions are currently being explored.

Algorithms for depth-related measures are often more complex for data sets which are not in general position than for data sets in general position. For example, the boundaries of subsequent halfspace depth regions are always disjoint when the data are in general position, but this does not hold for nongeneral position. Preferably, algorithms should be able to handle both the general position case and the nongeneral position case directly. As a quick fix, algorithms which were made

for general position can also be applied in the other case if one first adds small random errors to the data points. For large data sets, this ‘dithering’ will have a limited effect on the results.

2.1 Bivariate algorithms

Table 2 gives an overview of algorithms, each of which has been implemented, to compute the depth in a given point θ in \mathbb{R}^2 . These algorithms are time-optimal, since the problem of computing these bivariate depths has an $\Omega(n \log n)$ lower bound ([ACG⁺02], [LS00b]).

The algorithms for halfspace and simplicial depth are based on the same technique. First, data points are radially sorted around θ . Then a line through θ is rotated. The depth is calculated by counting the number of points that are passed by the rotating line in a specific manner. The planar arrangement depth algorithm is easiest to visualize in the regression setting. To compute the depth of a hyperplane H_θ with coefficients θ , the data are first sorted along the x -axis. A vertical line L is then moved from left to right and each time a data point is passed, the number of points above and below H_θ on both sides of L is updated.

Table 2: Computing the depth of a bivariate point.

DEPTH	TIME COMPLEXITY	SOURCE
Tukey depth	$O(n \log n)$	[RR96]
Simplicial depth	$O(n \log n)$	[RR96]
Arrangement/regression depth	$O(n \log n)$	[RH99]

In general, computing a median is harder than computing the depth in a point, because typically there are many candidate points. For instance, for the bivariate simplicial median the currently best algorithm requires $O(n^4)$ time, whereas its corresponding depth needs only $O(n \log n)$. The simplicial median seems difficult to compute because there are $O(n^4)$ candidate points (namely, all intersections of lines passing through two data points) and the simplicial depth regions have irregular shapes, but of course a faster algorithm may yet be found.

Fortunately, in several important cases the median can be computed without computing the depth of individual points. A linear-time algorithm to compute a bivariate centerpoint was described in [JM94]. Table 3 gives an overview of algorithms to compute bivariate depth-based medians. For the bivariate Tukey median the lower bound $\Omega(n \log n)$ was proved in [LS00], and the currently fastest algorithm takes $O(n \log^3 n)$ time ([LS03]). The lower bound $\Omega(n \log n)$ also holds for the median of arrangement (regression) depth as shown by [LS03b]. Fast algorithms were devised by [LS03b] and [VMR⁺08].

Table 3: Computing the bivariate median.

MEDIAN	TIME COMPLEXITY	SOURCE
Tukey median	$O(n \log^3 n)$	[LS03]
Simplicial median	$O(n^4)$	[ALS ⁺ 03]
Oja median	$O(n \log^3 n)$	[ALS ⁺ 03]
Regression depth median	$O(n \log n)$	[LS03b]

The computation of bivariate halfspace depth regions has also been studied. The first algorithm [RR96b] required $O(n^2 \log n)$ time per depth region. An algorithm to compute all regions in $O(n^2)$ time is constructed and implemented in [MRR⁺03]. This algorithm thus also yields the Tukey median. It is based on the dual arrangement of lines where topological sweep is applied. A completely different approach is implemented in [KMV02]. They make direct use of the graphics hardware to approximate the depth regions of a set of points in $O(nW + W^3) + nCW^2/512$ time, where the pixel grid is of dimension $(2W + 1) \times (2W + 1)$. Recently, [BRS11] constructed an algorithm to update halfspace depth and its regions when points are added to the data set.

2.2 Algorithms in higher dimensions

The first algorithms to compute the halfspace and regression depth of a given point in \mathbb{R}^d with $d > 2$ were constructed in [RS98] and require $O(n^{d-1} \log n)$ time. The main idea was to use projections onto a lower-dimensional space. This reduces the problem to computing bivariate depths, for which the existing algorithms have optimal time complexity. In [BCI+08] theoretical output-sensitive algorithms for the halfspace depth are proposed. An interesting computational connection between halfspace depth and multivariate quantiles was provided in [HPS10] and [KM12]. More recently, [DM14] provided a generalized version of the algorithm of [RS98] together with C++ code. For the depth regions of halfspace depth in higher dimensions an algorithm was recently proposed in [LMM14].

For the computation of projection depth see [LZ14]. The simplicial depth of a point in \mathbb{R}^3 can be computed in $O(n^2)$ time, and in \mathbb{R}^4 the fastest algorithm needs $O(n^4)$ time [CO01]. For higher dimensions, no better algorithm is known than the straightforward $O(n^{d+1})$ method to compute all simplices.

When the number of data points and dimensions are such that the above algorithms become infeasible, one can resort to approximate algorithms. For halfspace depth such approximate algorithms were proposed in [RS98] and [CMW13]. An approximation to the Tukey median using steepest descent can be found in [SR00]. In [VRH⁺02] an algorithm is described to approximate the deepest regression fit in any dimension.

3 Some other statistical techniques benefitting from computational geometry

Computational geometry has provided fast and reliable algorithms for many other statistical techniques.

Linear regression is a frequently used statistical technique. The ordinary least squares regression, minimizing the sum of squares of the residuals, is easy to calculate, but produces unreliable results whenever one or more outliers are present in the data. Robust alternatives are often computationally intensive. We here give some examples of regression methods for which geometric or combinatorial algorithms have been constructed.

1. L^1 regression. This well-known alternative to least squares regression minimizes the sum of the absolute values of the residuals, and is robust to vertical outliers. Algorithms for L^1 regression may be found in, e.g., [YKI⁺88] and [PK97].
2. Least median of squares (LMS) regression ([Rou84]). This method minimizes the median of the squared residuals and has a breakdown value of $1/2$. To compute the bivariate LMS line, an $O(n^2)$ algorithm using topological sweep has been developed [ES90]. An approximation algorithm for the LMS line was constructed in [MN⁺97]. The recent algorithm of [BM14] uses mixed integer optimization.
3. Median slope regression ([The50], [Sen68]). This bivariate regression technique estimates the slope as the median of the slopes of all lines through two data points. An algorithm with optimal complexity $O(n \log n)$ is given in [BC98], and a more practical randomized algorithm in [DMN92].
4. Repeated median regression ([Sie82]). Median slope regression takes the median over all couples (d -tuples in general) of data points. Here, this median is replaced by d nested medians. For the bivariate repeated median regression line, [MMN98] provide an efficient randomized algorithm.

The aim of cluster analysis is to divide a data set into clusters of similar objects. Partitioning methods divide the data into k groups. Hierarchical methods construct a complete clustering tree, such that each cut of the tree gives a partition of the data set. A selection of clustering methods with accompanying algorithms is presented in [SHR97]. The general problem of partitioning a data set into groups such that the partition minimizes a given error function f is NP-hard. However, for some special cases efficient algorithms exist. For a small number of clusters in low dimensions, exact algorithms for partitioning methods can be constructed. Constructing clustering trees is also closely related to geometric problems (see e.g., [Epp97], [Epp98]).

4 Other surveys

All results not given an explicit reference above may be traced in these surveys.

- [Mos13]: A survey of multivariate data depth and its statistical applications.
- [Sha76]: An overview of the computational complexities of basic statistics problems like ranking, regression, and classification.
- [Sma90]: An overview of several multivariate medians and their basic properties.
- [ZS00]: A classification of multivariate data depths based on their statistical properties.

References

- [ACG⁺02] G. Aloupis, C. Cortés, F. Gómez, M. Soss, and G.T. Toussaint. Lower bounds for computing statistical depth. *Comput. Stat. Data Anal.*, 40:223–229, 2002.
- [ALS⁺03] G. Aloupis, S. Langerman, M. Soss, and G.T. Toussaint. Algorithms for bivariate medians and a Fermat-Torricelli problem for lines. *Comput. Geom. Theory Appl.*, 26:69–79, 2003.
- [ABE⁺00] N. Amenta, M. Bern, D. Eppstein, and S.-H. Teng. Regression depth and center points. *Discrete Comput. Geom.*, 23:305–323, 2000.
- [BH99] Z.-D. Bai and X. He. Asymptotic distributions of the maximal depth estimators for regression and multivariate location. *Ann. Statist.*, 27:1616–1637, 1999.
- [BM14] D. Bertsimas and R. Mazumder. Least quantile regression via modern optimization. *Ann. Statist.*, 42:2494–2525, 2014.
- [BCI⁺08] D. Bremner, D. Chen, J. Iacono, S. Langerman, and P. Morin. Output-sensitive algorithms for Tukey depth and related problems. *Stat. Comput.*, 18:259–266, 2008.
- [BC98] H. Brönnimann and B. Chazelle. Optimal slope selection via cuttings. *Comput. Geom.*, 10:23–29, 1998.
- [BRS11] M. Burr, E. Rafalin, and S.L. Souvaine. Dynamic maintenance of halfspace depth for points and contours. arXiv:1109.1517 (2011).
- [Cha96] P. Chaudhuri. On a geometric notion of quantiles for multivariate data. *J. Amer. Statist. Assoc.*, 91:862–872, 1996.

- [CMW13] D. Chen, P. Morin, and U. Wagner. Absolute approximation of Tukey depth: theory and experiments. *Comput. Geom.*, 46:566–573, 2013.
- [CO01] A.Y. Cheng and M. Ouyang. On algorithms for simplicial depth. In *Proc. 13th Canadian Conf. on Comp. Geom.*, pages 53–56, Waterloo, 2001.
- [Cla14] G. Claeskens, M. Hubert, L. Slaets, and K. Vakili. Multivariate functional halfspace depth. *J. Amer. Statist. Assoc.*, 109:411–423, 2014.
- [DMN92] M.B. Dillencourt, D.M. Mount, and N.S. Netanyahu. A randomized algorithm for slope selection. *Internat. J. Comput. Geom. Appl.*, 2:1–27, 1992.
- [DG92] D.L. Donoho and M. Gasko. Breakdown properties of location estimates based on halfspace depth and projected outlyingness. *Ann. Statist.*, 20:1803–1827, 1992.
- [DM14] R. Dyckerhoff and P. Mozharovskyi. Exact computation of the halfspace depth. arXiv:1411.6927 [stat.CO], 2014.
- [ES90] H. Edelsbrunner and D.L. Souvaine. Computing least median of squares regression lines and guided topological sweep. *J. Amer. Statist. Assoc.*, 85:115–119, 1990.
- [Epp97] D. Eppstein. Faster construction of planar two-centers. In *Proc. 8th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 131–138, New Orleans, 1997.
- [Epp98] D. Eppstein. Fast hierarchical clustering and other applications of dynamic closest pairs. In *Proc. 9th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 619–628, San Francisco, 1998.
- [FFC12] H. Fritz, P. Filzmoser, and C. Croux. A comparison of algorithms for the multivariate L_1 -median. *Computation. Stat.*, 27:393–410, 2012.
- [Gow74] J.C. Gower. The mediancenter. *J. Roy. Statist. Soc. Ser. C*, 32:466–470, 1974.
- [HPS10] M. Hallin, D. Paindaveine, and M. Šiman. Multivariate quantiles and multiple-output regression quantiles: from L_1 -optimization to halfspace depth. *Ann. Statist.*, 38:635–669, 2010.
- [HRV01] M. Hubert, P.J. Rousseeuw, and S. Van Aelst. Similarities between location depth and regression depth. In L.T. Fernholz, editor, *Statistics in Genetics and in the Environmental Sciences*, pages 153–162. Birkhäuser Verlag, Basel, 2001.
- [HRS15a] M. Hubert, P.J. Rousseeuw, and P. Segaert. Multivariate functional outlier detection. *Stat. Methods & Applic.*, 24:177–202, 2015.

- [HRS15b] M. Hubert, P.J. Rousseeuw, and P. Segaert. Multivariate and functional classification using depth and distance. arXiv:1504.01128, 2015.
- [JM94] S. Jadhav and A. Mukhopadhyay. Computing a centerpoint of a finite planar set of points in linear time. *Discrete Comput. Geom.*, 12:291–312, 1994.
- [KM12] L. Kong and I. Mizera. Quantile tomography: using quantiles with multivariate data. *Statistica Sinica*, 22:1589–1610, 2012.
- [KMV02] S. Krishnan, N.H. Mustafa, and S. Venkatasubramanian. Hardware-assisted computation of depth contours. In *Proc. 13th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 558–567, 2002.
- [LS00] S. Langerman and W. Steiger. Computing a maximal depth point in the plane. In *Proc. Japan Conf. on Discrete and Computational Geometry*, pages 46–47, 2000.
- [LS00b] S. Langerman and W. Steiger. An optimal algorithm for hyperplane depth in the plane. In *Proc. 11th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 54–59, San Francisco, 2000.
- [LS03] S. Langerman and W. Steiger. Optimization in arrangements. *Lecture Notes in Computer Science 2607, Springer-Verlag, Berlin*, 2003.
- [LS03b] S. Langerman and W. Steiger. The complexity of hyperplane depth in the plane. *Discrete Comput. Geom.*, 30:299–309, 2003.
- [LCL12] J. Li, J.A. Cuesta-Albertos, and R.Y. Liu. DD-classifier: nonparametric classification procedure based on DD-plot. *J. Amer. Statist. Assoc.*, 107:737–753, 2012.
- [Liu90] R.Y. Liu. On a notion of data depth based on random simplices. *Ann. Statist.*, 18:405–414, 1990.
- [LPS99] R.Y. Liu, J. Parelius, and K. Singh. Multivariate analysis by data depth: descriptive statistics, graphics and inference. *Ann. Statist.*, 27:783–840, 1999.
- [LMM14] X. Liu, K. Mosler, and P. Mozharovskyi. Fast computation of Tukey trimmed regions in dimension $p > 2$. arXiv:1412.5122 [stat.CO], 2014.
- [LZ14] X. Liu and Y. Zuo. Computing projection depth and its associated estimators. *Stat. Comput.*, 24:51–63, 2014.
- [MMN98] J. Matoušek, D.M. Mount, and N.S. Netanyahu. Efficient randomized algorithms for the repeated median line estimator. *Algorithmica*, 20:136–150, 1998.

- [MRR⁺03] K. Miller, S. Ramaswami, P. Rousseeuw, J.A. Sellarès, D. Souvaine, I. Streinu, and A. Struyf. Efficient Computation of Location Depth Contours by Methods of Computational Geometry. *Stat. Comput.*, 13:153–162, 2003.
- [Miz02] I. Mizera. On depth and deep points: a calculus. *Ann. Statist.* 30:1681–1736, 2002.
- [Mos13] K. Mosler. Depth statistics. In C. Becker, R. Fried, and S. Kuhnt, editors, *Robustness and Complex Data Structures*, pages 17–34, Springer-Verlag, Berlin Heidelberg, 2013.
- [MN94] D.M. Mount and N.S. Netanyahu. Computationally efficient algorithms for high-dimensional robust estimators. *Graphical Models Image Proc.*, 56:289–303, 1994.
- [MN⁺97] D.M. Mount, N.S. Netanyahu, K. Romanik, R. Silverman, and A.Y. Wu. A practical approximation algorithm for the LMS line estimator. In *Proc. 8th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 473–482, New Orleans, 1997.
- [Oja83] H. Oja. Descriptive statistics for multivariate distributions. *Statist. Probab. Lett.*, 1:327–332, 1983.
- [PK97] S. Portnoy and R. Koenker. The Gaussian hare and the Laplacian tortoise: computability of squared-error versus absolute-error estimators. *Stat. Sci.*, 12:279–300, 1997.
- [Rou84] P.J. Rousseeuw. Least median of squares regression. *J. Amer. Statist. Assoc.*, 79:871–880, 1984.
- [RH99] P.J. Rousseeuw and M. Hubert. Regression depth. *J. Amer. Statist. Assoc.*, 94:388–402, 1999.
- [RH99b] P.J. Rousseeuw and M. Hubert. Depth in an arrangement of hyperplanes. *Discrete Comput. Geom.*, 22:167–176, 1999.
- [RR96] P.J. Rousseeuw and I. Ruts. Algorithm AS 307: Bivariate location depth. *J. Roy. Statist. Soc. Ser. C*, 45:516–526, 1996.
- [RR98] P.J. Rousseeuw and I. Ruts. Constructing the bivariate Tukey median. *Statistica Sinica*, 8:827–839, 1998.
- [RRT99] P.J. Rousseeuw, I. Ruts, and J.W. Tukey. The bagplot: A bivariate boxplot. *Amer. Statist.*, 53:382–387, 1999.
- [RS98] P.J. Rousseeuw and A. Struyf. Computing location depth and regression depth in higher dimensions. *Stat. Comput.*, 8:193–203, 1998.

- [RS04] P.J. Rousseeuw and A. Struyf. Characterizing angular symmetry and regression symmetry. *J. Statist. Plann. Inference*, 122:161–173, 2004.
- [RR96b] I. Ruts and P.J. Rousseeuw. Computing depth contours of bivariate point clouds. *Comput. Stat. Data Anal.*, 23:153–168, 1996.
- [Sen68] P.K. Sen. Estimates of the regression coefficient based on Kendall’s tau. *J. Amer. Statist. Assoc.*, 63:1379–1389, 1968.
- [Ser02] R. Serfling. A depth function and a scale curve based on spatial quantiles. In Y. Dodge, editor, *Statistical Data Analysis Based on the L1-Norm and Related Methods*, pages 25–38, Birkhäuser, Basel, 2002.
- [Sha76] M.I. Shamos. Geometry and statistics: problems at the interface. In J.F. Traub, editor, *Algorithms and Complexity: New Directions and Recent Results*, pages 251–280. Academic Press, Boston, 1976.
- [Sie82] A. Siegel. Robust regression using repeated medians. *Biometrika*, 69:242–244, 1982.
- [Sma90] C.G. Small. A survey of multidimensional medians. *Internat. Statistical Review*, 58:263–277, 1990.
- [SHR97] A. Struyf, M. Hubert, and P.J. Rousseeuw. Integrating robust clustering techniques in S-PLUS. *Comput. Stat. Data Anal.*, 26:17–37, 1997.
- [SR00] A. Struyf and P.J. Rousseeuw. High-dimensional computation of the deepest location. *Comput. Stat. Data Anal.*, 34:415–426, 2000.
- [The50] H. Theil. A rank-invariant method of linear and polynomial regression analysis (parts 1-3). *Nederl. Akad. Wetensch. Ser. A*, 53:386–392, 521–525, 1397–1412, 1950.
- [Tuk75] J.W. Tukey. Mathematics and the picturing of data. In *Proc. Internat. Congr. of Math.*, 2, pages 523–531, Vancouver, 1975.
- [VMR⁺08] M. van Kreveld, J. Mitchell, P. Rousseeuw, M. Sharir, J. Snoeyink, B. Speckmann. Efficient algorithms for maximum regression depth. *Discrete Comput. Geom.*, 39:656–677, 2008.
- [VRH⁺02] S. Van Aelst, P.J. Rousseeuw, M. Hubert, and A. Struyf. The deepest regression method. *J. Multivar. Anal.*, 81:138–166, 2002.
- [YKI⁺88] P. Yamamoto, K. Kato, K. Imai, and H. Imai. Algorithms for vertical and orthogonal L^1 linear approximation of points. In *Proc. 4th Sympos. Comput. Geom.*, pages 352–361, 1988.

[Zuo03] Y. Zuo. Projection based depth functions and associated medians. *Ann. Statist.*, 31:1460–1490, 2003.

[ZS00] Y. Zuo and R. Serfling. General notions of statistical depth function. *Ann. Statist.*, 28:461–482, 2000.